

Advances in Modeling the Generation of the Geomagnetic Field by the Use of Massively Parallel Computers and Profound Optimization

The Geomagnetic Field

At the Earth's surface, the magnetic field that is observed is similar to that that would be generated by a simple bar magnet running through the Earth's axis.

- This idea (permanent magnetism) was commonly believed a century ago.
- Because the temperature of the core is so high, permanent magnetism is not possible.
- Therefore, the magnetic field should decay, over tens of thousands of years.
- Since it does not, the field must be regenerating.
- Since the turn of the century, the idea that the core is molten iron which by moving generates a magnetic field arose.
- The set of equations to describe this are extremely non-linear and complex.
- Only in the last five to ten years have computers been able to solve these equations.



Thomas Clune, Daniel S. Katz, Gary A. Glatzmaier



Advances in Modeling the Generation of the Geomagnetic Field by the Use of Massively Parallel Computers and Profound Optimization

Geomagnetic Equations

$$\rho = \left[\left(\frac{\partial \rho}{\partial S} \right)_{\xi, p} S + \left(\frac{\partial \rho}{\partial \xi} \right)_{S, p} \xi + \left(\frac{\partial \rho}{\partial p} \right)_{S, \xi} p \right]$$

Equation of State

$$\frac{\partial \mathbf{B}}{\partial t} = \nabla \times (\mathbf{v} \times \mathbf{B}) - \nabla \times (\bar{\eta} \nabla \times \mathbf{B})$$

Magnetic Induction Equation

$$\nabla \cdot \bar{\rho} \mathbf{v} = 0$$

Mass Conservation

$$\nabla \cdot \mathbf{B} = 0$$

Magnetic Flux Conservation

$$\nabla^2 U = 4\pi G \rho$$

Gravitational Potential

$$\begin{aligned} \bar{\rho} \frac{\partial \mathbf{v}}{\partial t} = & -\nabla \cdot (\bar{\rho} \mathbf{v} \mathbf{v}) - \bar{\rho} \nabla (p / \bar{\rho} + U) - \left[\left(\frac{\partial \rho}{\partial S} \right)_{\xi, p} S + \left(\frac{\partial \rho}{\partial \xi} \right)_{S, p} \xi \right] \bar{\mathbf{g}} \hat{\mathbf{r}} \\ & + 2\bar{\rho} \mathbf{v} \times \boldsymbol{\Omega} + \nabla \cdot \left(2\bar{\rho} \bar{\mathbf{v}} \left(\bar{\mathbf{e}} - \frac{1}{3} (\nabla \cdot \mathbf{v}) \bar{\boldsymbol{\delta}} \right) \right) + \frac{1}{\mu_0} (\nabla \times \mathbf{B}) \times \mathbf{B} \end{aligned}$$

Momentum Equation

$$\begin{aligned} \bar{\rho} \frac{\partial S}{\partial t} = & -\nabla \cdot (\bar{\rho} S \mathbf{v}) + \nabla \cdot (\bar{\rho} \bar{\kappa} \nabla S) + \frac{1}{\bar{T}} \frac{d}{dr} \left(r^2 c_p \bar{\rho} \bar{\kappa}^{-1} \frac{dT}{dr} \right) \\ & + \frac{\bar{\eta}}{\mu_0 \bar{T}} |\nabla \times \mathbf{B}|^2 + \frac{\bar{\mathbf{g}}}{\bar{T}} \left[\bar{\kappa} \left(\frac{\partial \rho}{\partial S} \right)_{\xi, p} \frac{\partial S}{\partial r} + \bar{\kappa}^{-\xi} \left(\frac{\partial \rho}{\partial \xi} \right)_{S, p} \frac{\partial \xi}{\partial r} \right] \end{aligned}$$

Heat Equation

$$\bar{\rho} \frac{\partial \xi}{\partial t} = -\nabla \cdot (\bar{\rho} \xi \mathbf{v}) + \nabla \cdot (\bar{\rho} \bar{\kappa}^{-\xi} \nabla \xi)$$

Compositional Equation



Thomas Clune, Daniel S. Katz, Gary A. Glatzmaier



Advances in Modeling the Generation of the Geomagnetic Field by the Use of Massively Parallel Computers and Profound Optimization

Thomas Clune

Silicon Graphics, Inc.

Daniel S. Katz

Jet Propulsion Laboratory

California Institute of Technology

Gary A. Glatzmaier

University of California, Santa Cruz



**PARALLEL PROCESSING
FOR SCIENTIFIC COMPUTING**




Advances in Modeling the Generation of the Geomagnetic Field by the Use of Massively Parallel Computers and Profound Optimization

Magnetic Field Lines

Blue lines are directed inward

Red lines are directed outward



(Graphic generated by DYNAMO)



Thomas Clune, Daniel S. Katz, Gary A. Glatzmaier



Advances in Modeling the Generation of the Geomagnetic Field by the Use of Massively Parallel Computers and Profound Optimization

DYNAMO Code

Development from 1981 to present by Gary Glatzmaier

- Based initially on a solar convection code

Spectral transform method

- All variables are expanded in spherical harmonics for horizontal structure
- Causes decoupling of spherical degree and order components of variables
- Avoids the pole problem
- Products are computed in grid space

Using Fourier and Legendre transforms

For small problems, this leads to high accuracy with low cost

For large problems, Legendre transforms are an increasing part of the work

- On parallel machines, they can be done very quickly

Historically, spectral codes were expected to be poor performers on MPPs

- Global communication pattern requires every node to communicate with every other node at every time step



Thomas Clune, Daniel S. Katz, Gary A. Glatzmaier



Advances in Modeling the Generation of the Geomagnetic Field by the Use of Massively Parallel Computers and Profound Optimization

Results of DYNAMO runs

Fields have correct amplitude and structure outside the core

- Agree with measurements on Earth's surface and in space

Time dependence on the short and long time scales agree with measurements

- Short time scale: drift of non-dipolar part is westward at correct speed
- Long time scale: magnetic field reversal as seen in paleomagnetic record

Predicted eastward rotation of inner core with respect to the Earth's surface

- Seismic analyses now show support for this rotation



Thomas Clune, Daniel S. Katz, Gary A. Glatzmaier



Advances in Modeling the Generation of the Geomagnetic Field by the Use of Massively Parallel Computers and Profound Optimization

Serial Computational Issues

As problems grow (increasing the spatial resolution) as $O(n)$, the work in the Legendre transforms grows as $O(n^4)$

- This is the first place to look for optimizations
- We observe that Legendre transforms can be written as matrix-matrix multiplications, because they are independent of r (radial level)
- This implies that level 3 BLAS can be used

These are vendor optimized

On most machines, these produce the best performance

- Further observation: matrix-matrix multiplication can be decomposed into the sum and difference of two smaller matrix-matrix multiplications

Savings is almost a factor of 2

Perform transforms for both the Northern and Southern hemispheres simultaneously with the same number of operations required for a single hemisphere.



Thomas Clune, Daniel S. Katz, Gary A. Glatzmaier



Advances in Modeling the Generation of the Geomagnetic Field by the Use of Massively Parallel Computers and Profound Optimization

More Serial Optimization Issues

Azimuthal transform is FFT, while radial is Chebyshev transforms

- Both can use vendor-supplied FFT libraries

Good performance is likely

Other operations can also be replaced by vendor-supplied routines

- In DYNAMO, LU decomposition is an example

Ensuring cache friendly coding

- Use stride 1 references
- Reorder loops so that inner loop varies most quickly
- DYNAMO did not require many changes

Instead of transforming the entire set of fields to grid space at once, only small subset of the spatial domain are computed at any one time.

- This is both cache- and memory-friendly



Thomas Clune, Daniel S. Katz, Gary A. Glatzmaier



Advances in Modeling the Generation of the Geomagnetic Field by the Use of Massively Parallel Computers and Profound Optimization

Parallel Implementation

Major operations are done locally on each processor, using global transposes to rearrange data as needed

- Legendre transforms (co-latitude and spherical harmonic degree are local)
- Chebyshev transforms and implicit solves (radius and Chebyshev degree are local)
- Fourier transform (azimuthal wave number and longitude are local)

The two remaining dimensions are both decomposed over the processors

- This allows more processors to be used on any given problem size than if the decomposition was just over one remaining dimension

For the implicit solves, the method of decomposing the two remaining dimensions (spherical harmonic degree and azimuthal wave number) is very important for performance

- A judicious arrangement of data distribution exists which allows good cache reuse (the matrices can operate on several RHS simultaneously)




Thomas Clune, Daniel S. Katz, Gary A. Glatzmaier



Advances in Modeling the Generation of the Geomagnetic Field by the Use of Massively Parallel Computers and Profound Optimization

Data Distribution for Implicit Solves



Data layout used in the spectral domain for Legendre transposes and accumulating linear terms

Data layout used for matrix solves during implicit step

In addition to performance, this distribution scheme saves memory, since the matrices to be solved depend on l .



Thomas Clune, Daniel S. Katz, Gary A. Glatzmaier



Advances in Modeling the Generation of the Geomagnetic Field by the Use of Massively Parallel Computers and Profound Optimization

Cray T3E-Specific Optimizations

E registers

- Hardware mechanism that allows processors to directly access memory on remote processors, bypassing the cache
- Can be used for non-unit stride in local memory
- Allows very efficient local transposes

Shmem libraries

- Low latency, high bandwidth communications library on T3E
- Requires one-sided calls
- DYNAMO is careful to use long, unit-stride messages
- DYNAMO is careful to schedule communications to avoid hotspots on the network

Streams

- Cache pre-fetching mechanism for uniform memory references



Thomas Clune, Daniel S. Katz, Gary A. Glatzmaier



Advances in Modeling the Generation of the Geomagnetic Field by the Use of Massively Parallel Computers and Profound Optimization

Performance

The original code was sequential and could be run on 4 processors of a Cray T90

- Largest problem that was run: $65 \times 150 \times 300$ (radius \times latitude \times longitude)
- Overall performance: 1.0 GFLOP/s
- Wall-clock time to solve this problem for 100 years: 24 hours

The new parallel code has been run on 1488 T3E-1200 processors

- The largest problem that has been run: $193 \times 1858 \times 1488$
- Overall performance: 637 GFLOP/s
- The wall-clock time to solve this problem for 100 years: 1 year

Performance improvement: 600 times

Wall-clock improvement: 3 times (for the small original problem size)

(The current parallel code solves the Coriolis term explicitly. This dramatically reduces the maximum time step. This can be fixed, and should increase the maximum time step by 50 times.)

Most importantly, the maximum possible problem size has increased by 180 times




Thomas Clune, Daniel S. Katz, Gary A. Glatzmaier



Advances in Modeling the Generation of the Geomagnetic Field by the Use of Massively Parallel Computers and Profound Optimization

Scaled-Size Problems



The problem size is a function of $NN*NI*NJ$

The work is a function of $NN*NI*NI*NH$




Thomas Clune, Daniel S. Katz, Gary A. Glatzmaier



Advances in Modeling the Generation of the Geomagnetic Field by the Use of Massively Parallel Computers and Profound Optimization

Fixed-Size Problems



Thomas Clune, Daniel S. Katz, Gary A. Glatzmaier



Advances in Modeling the Generation of the Geomagnetic Field by the Use of Massively Parallel Computers and Profound Optimization

Future Work

Add capability to perform run-time analyses

- Summarize / Average sections / domains of fields
- Generate animations on-the-fly

Add MPI switches

- First, MPI-1 two-sided calls
- Then, MPI-2 one-sided calls

Port to other machines

- Origin 2000 is the first target
- Then clusters, SPs, others

Cleaning-up and documenting the code

Pre- and post-processing tools are becoming bottlenecks

- This is true in many (most?) parallel disciplines



Thomas Clune, Daniel S. Katz, Gary A. Glatzmaier



Advances in Modeling the Generation of the Geomagnetic Field by the Use of Massively Parallel Computers and Profound Optimization

Conclusions

On tightly coupled machines, such as the T3E, spectral transform codes work well

- As problem size increases, communications time decreases

Previous work has been done at lower resolutions

- Solving for large-scale features has been done by making approximations to eliminate the small-scale details
- This requires material parameters which are known to be non-physical, but do allow computational solution

High resolution studies will enable new science

- Smaller-scale features appear
- Realistic material parameters can be used
- An understanding of the required resolution will be obtained, which would provide much more confidence in computational solutions



Thomas Clune, Daniel S. Katz, Gary A. Glatzmaier

